

****Volume Title****

*ASP Conference Series, Vol. **Volume Number***

****Author****

© ****Copyright Year**** *Astronomical Society of the Pacific*

Map-making for large-format detector arrays on CCAT

Gaelen Marsden¹, Tim Jenness², and Douglas Scott¹

¹*Department of Physics and Astronomy, University of British Columbia,
6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada*

²*Department of Astronomy, Cornell University, Ithaca NY, 14853, USA*

Abstract. CCAT is a large submillimetre telescope to be built near the ALMA site in northern Chile. A large-format KID camera, with up to 48,000 detectors at a single waveband sampled at ~ 1 kHz, will have a data rate ~ 50 times larger than SCUBA-2, the largest existing submillimetre camera. Creating a map from this volume of data will be a challenge, both in terms of memory and processing time required. We investigate how to extend SMURF, the iterative map-maker used for reducing SCUBA-2 observations, to a distributed-node parallel system, and estimate how the processing time scales with the number of nodes in the system.

1. Introduction

CCAT (see e.g. Jenness et al. 2014) will be a large telescope optimized for observations in the submillimetre and will be located near the ALMA site in northern Chile. One of the proposed instruments is SWCam (Stacey et al. 2014), a large-format kinetic inductance detector (KID) array with subarrays spanning $850\text{--}350\,\mu\text{m}$, and possibly extending to $200\,\mu\text{m}$. With a field of view of 16 arcmin and a sampling rate of ~ 1 kHz, SWCam will be able to quickly survey large fields (tens of square degrees) at high resolution ($3.5''$ at $350\,\mu\text{m}$) and sensitivity; it will be able to reach noise levels five times below the confusion limit at each wavelength over 35 deg^2 in the first year of observations.

Current submillimetre photometric instruments, such as SHARC-II (Dowell et al. 2003), LABOCA (Siringo et al. 2009), ArTéMiS (Revéret et al. 2014) and SCUBA-2 (Holland et al. 2013), have hundreds to thousands of detectors sampled at tens to hundreds of Hertz. The SWCam data volume of 5 TB/night will thus be considerably larger than current data sets, and we must ensure that map-making software will be able to reduce the data in a reasonable amount of time. Marsden et al. (2014) discuss the problem of scaling current map-makers to SWCam, and suggest distributed-memory parallelization as a solution to handling the large data sets. Here, we further develop the distributed-memory parallel model and present results of timing tests.

2. Distributed Memory Parallel Map-making

Following Marsden et al. (2014), we focus on SMURF (Chapin et al. 2013, ascl:1310.007), the iterative map-maker used for reduction of SCUBA-2 data. A simplified flow chart

describing the algorithm is shown in Figure 1. Very briefly, the data are read from disk,




Figure 1. Flow chart illustrating iterative map-making using N models. The models represent sources of signal and noise, such as common-mode signal, low frequency $1/f$ noise, and astronomical signal.

preprocessing steps (e.g. spike detection) are performed, then a number of noise/signal models are fit and removed from the data.¹ At the end of the sequence, the r.m.s. of the time stream residuals is compared to the previous iteration; if the change is larger than the convergence tolerance, the sequence of models is re-fit to the new residuals. The iteration loop continues until the convergence condition is met, up to a maximum number of iterations.

Marsden et al. (2014) suggest that, even accounting for advances in computer hardware in the next five years, a single machine will not be able reduce more than 15 minutes of data at a time, and will not be able to keep up with data collection. Adapting the iterative algorithm to run on a distributed-memory parallel computer would allow for the reduction of data sets longer than 15 minutes (necessary for making high-fidelity maps) and for the reduction run time. Figure 2 shows how this can be accomplished. For M processors, we divide the data set into M chunks, with each processor handling




Figure 2. Distributed-memory parallel version of the iterative map-maker. Each process is responsible for a distinct chunk of data. Communication between processors is necessary for most (but not all) models. The two-processor version is shown—this is easily extended to M processors. Hierarchical communication between all processors is required.

only its chunk of data. Communication between the processes is required when calculating models that depend on the full data set. This communication is represented by the

¹The number, type and sequence of models fit is a configuration option.

large green arrows in Figure 2. Other than these communication steps, the processors run independently.

For the most part, only (relatively) small data arrays need to be shared. As an example, the common mode model (a time-varying spatially-independent signal seen by all detectors), each process accumulates the sum of signal over detectors in its data chunk, then the accumulated sums are combined over all processes using collective communication. For the astronomical sky model (in other words, the sky map), each process accumulates signal into its own copy of the sky map and the maps are shared in the same manner as for the common mode model. For models such as the high-pass filter (which reduces low-frequency detector noise and atmospheric noise), no communication is necessary, since the result for each detector is independent of all others.

3. Prototype and Timing Tests

We have written a prototype map-maker² to test the parallel iterative algorithm described above. The code is written in C and uses `Open MPI`.³ The prototype implements the three models discussed: common mode, high-pass filter and astronomical signal.

The prototype has been run on a range of data set sizes and a varying number of processors in order to determine how the run time scales with the number of processors. The results of these tests are shown in Figure 3. The scaling of run times with number of processors for each data set size is well-described by a power law, and all three are reasonably well-fit by a single power-law index (shown in the figure as orange dotted lines), $\alpha = -0.762$. This corresponds to a reduction of run time by a factor of 1.7 for each factor of two in the number of processors, or a factor of ~ 20 for 50 processors. This is sufficient to reduce the run time of SWCam data to better than rate at which the data are collected.

4. Conclusions

We have demonstrated that the iterative map-maker used by SCUBA-2 can be adapted for use on a distributed-memory parallel system. A relatively modest cluster of 50 processors will be able to reduce SWCam data faster than they are collected, and reduces the memory required per processor to ~ 200 GB per hour of data. We are now adapting SMURF so that it can be run on such a system.

Acknowledgments. Thanks to Jack Sayers for providing the data simulation code that was used as the input to the timing tests. The SWCam prototype map-maker makes use of facilities provided by WestGrid and Compute Canada Calcul Canada.⁴

References

Chapin, E. L., Berry, D. S., Gibb, A. G., Jenness, T., Scott, D., Tilanus, R. P. J., Economou, F., & Holland, W. S. 2013, MNRAS, 430, 2545. arXiv:1301.3652

²Available at: <https://github.com/CCAT0bservatory/mpi-mapmaker-test>.

³An open source implementation of the Message Passing Interface. See: <http://www.open-mpi.org/>.

⁴<http://www.computecanada.ca>




Figure 3. Timing tests for the prototype parallel iterative map-maker. Run time vs. number of processors for three simulated data sets is shown. Each data set is 30 minutes of data sampled at 1500 Hz for the indicated number of detectors. The run times show power-law scaling. A power-law fit to each data set is shown, as is a fit to all three data sets with a common power law index (orange dotted curves).

- Dowell, C. D., et al. 2003, in *Millimeter and Submillimeter Detectors for Astronomy*, edited by T. G. Phillips, & J. Zmuidzinas, vol. 4855 of *Proc. SPIE*, 73
- Holland, W. S., et al. 2013, *MNRAS*, 430, 2513. [arXiv:1301.3650](#)
- Jenness, T., et al. 2014, in *Software and Cyberinfrastructure for Astronomy III*, edited by G. Chiozzi, & N. M. Radziwill, vol. 9152 of *Proc. SPIE*, 91522W. [arXiv:1406.1515](#)
- Marsden, G., Brazier, A., Jenness, T., Sayers, J., & Scott, D. 2014, in *Astronomical Data Analysis Software and Systems XXIII*, edited by N. Manset, & P. Forshay, vol. 485 of *ASP Conf. Ser.*, 399. [arXiv:1405.0482](#)
- Revéret, V., et al. 2014, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, edited by W. S. Holland, & J. Zmuidzinas, vol. 9153 of *Proc. SPIE*, 915305. [arXiv:1407.2545](#)
- Siringo, G., et al. 2009, *A&A*, 497, 945. [arXiv:0903.1354](#)
- Stacey, G., et al. 2014, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, edited by W. S. Holland, & J. Zmuidzinas, vol. 9153 of *Proc. SPIE*, 91530L